LOCAL CORRELATION-BASED CIRCUITRY CAN ACCOUNT FOR RESPONSES TO MULTI-GRATING STIMULI IN A MODEL OF CAT V1.

T.Z. Lauritzen[1,4], A.E. Krukowski[1,4,*], and K.D. Miller[1−4]
[1]Graduate Group in Biophysics
[2]Depts. of Physiology and Otolaryngology
[3]Sloan-Swartz Center for Theoretical Neurobiology
[4]W.M. Keck Center for Integrative Neuroscience
University of California
San Francisco, CA 94143-0444

[*]Current address for AEK: NASA Ames Research Center, Mail Stop 262-2, Moffett Field, CA 94035-1000

Abbreviated title: Model of multi-grating suppression in cat V1

Address Correspondence To:
Kenneth D. Miller
Dept. of Physiology
University of California
513 Parnassus
San Francisco, CA 94143-0444

Phone: 415-476-8217
Fax: 415-476-4929
Electronic Mail: ken@phy.ucsf.edu

Abstract

In cortical simple cells of cat striate cortex, the response to a visual stimulus of the preferred orientation is partially suppressed by simultaneous presentation of a stimulus at the orthogonal orientation, an effect known as "cross-orientation inhibition". It has been argued that this is due to the presence of inhibitory connections between cells tuned for different orientations, but intracellular studies suggest that simple cells receive inhibitory input primarily from cells with similar orientation tuning. Furthermore, response suppression can be elicited by a variety of non-preferred stimuli at all orientations. Here we study a model circuit that was presented previously to address many aspects of simple cell orientation tuning, which is based on local intracortical connectivity between cells of similar orientation tuning. We show that this model circuit can account for many aspects of cross-orientation inhibition and, more generally, of response suppression by non-preferred stimuli and of other non-linear properties of responses to stimulation with multiple gratings.

Introduction

Cells in cat primary visual cortex (V1) are tuned for the orientation of light/dark borders (Hubel and Wiesel 1962). Understanding the circuity underlying this orientation selectivity remains a central problem in systems neuroscience (reviewed in Ferster and Miller 2000).

Clues to the underlying circuitry are provided by experiments involving the superposition of two drifting sinusoidal luminance gratings. A typical simple cell in layer 4 of cat V1 responds to a drifting grating shown at its preferred orientation and is silent in response to a drifting grating of the perpendicular (null) orientation. However, superposition of the null grating with the preferred causes a reduction in response relative to the response to the preferred grating alone, a phenomenon known as "cross-orientation inhibition" (Bonds 1989; DeAngelis et al. 1992; Morrone et al. 1982). More generally, superposition of a non-preferred grating can suppress responses to a preferred grating (Bonds 1989; DeAngelis et al. 1992).

This suppression suggests that a non-preferred grating recruits inhibition and/or disrupts the cell's excitatory drive. In particular, the existence of cross-orientation inhibition has led to the suggestion that inhibition from cells tuned to dissimilar orientations plays an important role in setting the gain of cortical responses (Carandini and Heeger 1994; Carandini et al. 1997, 1999; Heeger 1992; Heeger et al. 1996). However, evidence from intracellular recordings argues against this idea. Such recordings show that the excitation and the inhibition received by simple cells in cat layer 4 show similar orientation tuning, with both peaked at the preferred orientation and falling to small values at the orthogonal orientation (Anderson et al. 2000a; Ferster 1986), and that the orientation selectivity of voltage responses is neither created, nor sharpened, by intracortical circuitry (Chung and Ferster 1998; Ferster et al. 1996).

In this paper, we show that a simple model circuit that is consistent with the intracellular data can account for many of the two-grating suppression effects, including cross-orientation inhibition. This model circuit was originally inspired by the findings that the inhibition and excitation received by a layer 4 simple cell have similar orientation tuning (Anderson et al. 2000a; Ferster 1986), but are in a "push-pull" or spatially opponent relationship (Ferster

1988; Hirsch et al. 1998): in ON-subregions, where light evokes excitation, dark evokes inhibition, and similarly dark evokes excitation and light evokes inhibition in OFF-subregions. Accordingly, we proposed (Troyer et al. 1998) that excitatory cells tend to make connections onto cells of similar preferred orientation and similar absolute spatial phase (similar locations in visual space of ON-subregions and of OFF-subregions), while inhibitory cells tend to project to cells of similar preferred orientation and opposite absolute spatial phase. In addition, we assumed that the LGN connections to a simple cell are organized in an oriented, subregion-specific manner (Chung and Ferster 1998; Ferster et al. 1996; Hubel and Wiesel 1962; Reid and Alonso 1995; Tanaka 1983): ON-center LGN inputs have receptive fields aligned over the simple cell's ON subregions, and OFF-center inputs are aligned on OFF subregions. We showed (Troyer et al. 1998) that, provided that the LGN-driven inhibition was stronger than the direct LGN excitation, this circuitry could account for the invariance with stimulus contrast of orientation tuning (Sclar and Freeman 1982; Skottun et al. 1987) and for a number of other intracellular and extracellular observations. Here we show that this circuitry can also account for cross-orientation inhibition, and more generally for a variety of two-grating suppression phenomena.

One of the more extensive experimental studies of such suppression was done by Bonds (1989). He used as a visual stimulus a preferred (base) grating at one temporal frequency and a superposed mask grating at a different temporal frequency. By analyzing the two corresponding temporal components of the response, he separated the base-driven and mask-driven response components, and studied how each component depended on the orientation, contrast, and spatial frequency of the mask grating. In our modeling studies we follow this procedure to determine how closely our model is able to reproduce these experimental observations. We also address two other experiments examining responses to superpositions of gratings with different temporal frequencies at the preferred orientation (Dean et al. 1982; Reid et al. 1992). These experiments found that the modulation of the cell response to low-temporal-frequency stimuli is suppressed by the superposition of high-temporal-frequency stimuli, while the modulation of the cell response to high-temporal-frequency stimuli is en-

hanced by superposition of low-temporal-frequency stimuli. Here we reproduce these results and explain some of the findings by a toy model.

Methods

We study a model previously described ("computational model" of Troyer et al. 1998). The major difference in the present work is that N-methyl-D-aspartate (NMDA) receptors have been included at excitatory synapses onto excitatory cells. Here we present only the basics of that model along with details of any differences in the present implementation.

Our LGN model consists of 7200 LGN X-cells arranged in four overlying 30×30 sheets of ON cells and four similar sheets of OFF cells, with ON and OFF lattices offset by one-half square lattice spacing, covering $6.8 \times 6.8°$ of the visual field. LGN firing rates in response to a single grating were calculated by assuming rates were sinusoidally modulated, up to rectification at zero rate, about background rates of 10 and 15 Hz for ON and OFF cells respectively, with the amplitude of the sinusoidal modulation chosen so that the first harmonic of the rectified responses matched values at a given contrast and temporal frequency reported by Sclar (1987). To compute responses to multiple gratings, we added the sinusoidal rate modulations induced by each grating and then rectified the result. Spikes were then generated from these rates in a random (Poisson) fashion. Overlaying cells have 25% correlation in their spike trains (each of four overlaying cells picked spikes with probability 1/4 from a common set of four Poisson spike trains), to match data showing correlations among LGN cells with overlapping receptive fields (Alonso et al. 1996).

The cortical model includes 1600 excitatory- and 400 inhibitory layer 4 simple cells, representing a $2/3 \times 2/3$ mm patch of cortex, corresponding to $0.75 \times 0.75°$ in visual angle. Each cortical cell was assigned a Gabor-function describing its receptive field, with orientation given by a measured orientation map from cat V1, spatial frequency of 0.8 cycles/degree, retinotopic position progressing uniformly across the sheet, and spatial phase assigned randomly to each cell. The precise parameters used for the Gabor functions were

those specifying the "broadly tuned" cells in Troyer et al. (1998), which were designed to reproduce the observed (Anderson et al. 2000a; Ferster et al. 1996) 35^o half-width at half-height of the orientation tuning of intracellular voltage modulations in response to optimal sinusoidal gratings. The connection strength from LGN cells to a cortical cell was determined as in Troyer et al. (1998) by a probabilistic sampling of the cell's Gabor function, where positive regions of the Gabor are converted to probabilities of a connection from an ON-cell centered on that point, and negative regions converted to probabilities of OFF-cells connecting.

Cortical cells were modeled as single-compartment conductance-based integrate-and-fire neurons as in Troyer et al. (1998), with the following differences. The cortical background excitatory input (Poisson) received by all cells was set to a rate of 6000 Hz, resulting in background firing rates of approximately 0.5 Hz for excitatory cells and 20-30 Hz for inhibitory cells. The amplitude of the adaptation conductance, $\bar{g}_{\mathrm{adapt}}$, was reduced by a factor of 5 from a value of 3 nS to 0.6 nS, to bring firing rates up to more realistic levels (as discussed in Troyer et al. 1998); this yielded reasonable levels of excitatory cell gain as measured from plots of firing rate versus instantaneous membrane potential.

Finally, synaptic conductances were identical to those used in Troyer et al. (1998) except that we have included NMDA-mediated conductances, in addition to the AMPA-mediated conductances, in all of the excitatory synapses except for the thalamocortical synapses onto inhibitory cells, which were purely AMPA mediated. NMDA receptors were included simply because we believe they make the model more realistic, and allow the model to explain cortical temporal frequency tuning (Krukowski and Miller 2001) as well as orientation tuning (Troyer et al. 1998); inclusion of NMDA receptors has therefore become the "default" for our lab's studies of the correlation-based circuit. We used parameters determined in other studies (Krukowski 2000; Krukowski and Miller 2001) and did not tune parameters to address the present issues. The inclusion of NMDA receptors impacts the present issues only in reducing the F1/DC ratio at higher temporal frequencies, which turns out to be crucial to explaining the effects of mixing low- and high-temporal frequency gratings (see sections titled

"Modulation changes with multiple gratings" in Results and Discussion).

The decay of the NMDA conductances is modeled as a double exponential with a fast and a slow time constant:

$$g_{\mathrm{NMDA}}(t) = \sum_{t_j < t} \overline{g}_{\mathrm{NMDA}}(V) \left(f_{fast} e^{-(t-t_j)/\tau^{fall}_{\mathrm{NMDA},fast}} + (1 - f_{fast}) e^{-(t-t_j)/\tau^{fall}_{\mathrm{NMDA},slow}} - e^{-(t-t_j)/\tau^{rise}_{\mathrm{NMDA}}} \right)$$

where the sum is over presynaptic spike times t_j, and f_{fast} represents the contribution of the faster exponential to the total decay term. Parameters were taken from data for adult rats in a developmental study of NMDA conductances in the rat visual cortex (Carmignoto and Vicini 1992): $\tau^{fall}_{\mathrm{NMDA,fast}} = 63$ msec, $\tau^{fall}_{\mathrm{NMDA,slow}} = 200$ msec, $f_{fast} = 88\%$. We chose $\tau^{rise}_{\mathrm{NMDA}} = 5.5$ msec to set the 10-90% rise time of the NMDA EPSC to be equal to 7.8 msec as has been observed experimentally (Lester et al. 1990). The voltage dependence of $\overline{g}_{\mathrm{NMDA}}$ followed the model described in Jahr and Stevens (1990). The relative strength of NMDA and AMPA conductances were set in terms of the integrated current (*i.e.*, the total charge transfer) through excitatory conductances when the postsynaptic cell is clamped at the spike-threshold voltage. 90% of the integrated current in thalamocortical synapses to excitatory cells was mediated by NMDA, which is the value obtained by matching AMPA and NMDA amplitudes to those observed at thalamocortical synapses at the oldest ages studied in thalamocortical slices (Crair and Malenka 1995), and 95% of the current in intracortical excitatory synapses was mediated by NMDA.

The probability that a cortical cell connected to any other cortical cell depended on the correlation between their sets of LGN inputs, as in Troyer et al. (1998): the probability of a connection from an excitatory cell monotonically increased with the degree of correlation, while the probability of a connection from an inhibitory cell monotonically increased with the degree of anticorrelation. This connectivity rule yields the basic cortical circuit structure within a single iso-orientation column shown in cartoon form in figure 1A. For simplicity, the inhibitory cells received only thalamocortical input. Previous simulations have demonstrated that intracortical excitatory connections onto inhibitory cells do not have a significant effect on the behavior of the model (Krukowski 2000; Troyer et al. 1998). An important feature of the model is that the inhibition is dominant over feedforward excitation, rather than

precisely balancing it.

Simulations for a given grating or multi-grating stimulus were run as follows. We first allowed the network to run for one second of simulated time while LGN cells fired at their background rates. We then started the grating stimulus and allowed the simulation to run an additional 0.25 seconds to suppress transients, then recorded data for 1 second of simulated time as the grating stimulus continued. PSTHs of 20 such simulations (each with different seeds for the random number generator controlling the Poisson sampling of LGN spikes) were made for the spiking results. Cells of all orientation preferences are represented in the cortical network. To be certain of avoiding artifacts of alignment of the grid with the stimulus, the preferred-orientation stimulus was always at an orientation of 128°, and we sampled our results from the 35 excitatory cells with preferred orientations within ±2.5° of 128°.

Toy rate model

In modeling responses to sums of two grating stimuli at the preferred orientation at different temporal frequencies, we considered both the full model as described above and a simple toy model. The toy model was based on the DC and F1 of the feedforward input (the LGN-driven input, both direct LGN excitation and indirect inhibition via interneurons) observed in the full model. We assumed that the feedforward inputs due to the two sinusoids were simply added and the result passed through a linear threshold input/output curve. We considered 2 sine curves: $y_{2Hz}(t) = \sin(4\pi t/sec) - \beta_{2Hz}$, $y_{8Hz}(t) = \alpha(\sin(16\pi t/sec) - \beta_{8Hz})$, where α is the amplitude of the 8 Hz grating relative to that of the 2 Hz grating, and the β's are the *inhibitory* total feedforward DC (the sum of LGN excitation and inhibition from LGN-driven inhibitory interneurons). We estimated the size of the β's in the full model by examining the total synaptic current evoked when the membrane potential is clamped at spike threshold. We considered response to a 2 Hz grating alone and an 8 Hz grating alone, where grating contrasts were chosen so that evoked current F1's were comparable (2 Hz grating at 20% contrast, 8 Hz grating at 40% contrast, ratio of 2 Hz F1 to 8 Hz F1 is

approximately 0.9). Defining β as DC/F1 of the resulting currents, we found $\beta_{8Hz} = 0.42$ and $\beta_{2Hz} = 0.019$. We considered a linear rectified model of response $r(t)$ to an input waveform $i(t)$: $r(t) = [i(t) - \theta]^+$, where θ is a threshold (values used given in Results) and $[x]^+ = x$, $x > 0$; $[x]^+ = 0$, $x \leq 0$. To model single sinusoids, $i(t)$ is set equal to $y_{2Hz}(t)$ or $y_{8Hz}(t)$; for double sinusoids, $i(t) = y_{2Hz}(t) + y_{8Hz}(t)$. We then examined the Fourier component of $r(t)$ at 2 Hz and/or 8 Hz.

Results

Basic concepts of the model circuit

To understand our results, it is important to recall the main ideas of our circuit model, which will figure prominently in the following results (Fig. 1). We define feedforward input *Place* to excitatory cells as the sum of the direct LGN excitatory input and the LGN-driven input *Fig. 1* from inhibitory neurons. Due to the dominant antiphase inhibition, the mean feedforward *about* input evoked by grating stimuli is inhibitory. Cortical excitatory simple cells can only be *here* driven to fire by the temporal modulation of this input: inhibition and excitation are driven at opposite phases of the modulation, so that one goes up when the other goes down. This modulation allows excitation to periodically dominate over inhibition and drive responses, even though inhibition dominates in the mean (Fig. 1B).

The mean LGN input evoked by a grating does not depend on grating orientation (because this mean is just the weighted sum of the mean rates of the individual LGN inputs to the simple cell, and the responses of LGN cells are assumed to be untuned for orientation). But the modulation of the LGN input varies strongly with orientation. In response to a preferred orientation stimulus, all of the LGN inputs to a cell modulate their firing rates together, so the total LGN input to the cell is strongly modulated and the cell periodically fires (Fig. 1B). As the stimulus orientation is moved away from the preferred, the different LGN inputs to a cell come to be increasingly desynchronized in their rate modulations. These modulations at different phases wash out, so that the net input to the cell becomes temporally steady

and unmodulated, albeit with the same mean, and the cell is inhibited (Fig. 1C).

The net feedforward input – that is, the sum of the LGN input and the antiphase inhibition – is determined by the LGN input as follows. The mean LGN input evokes a net inhibitory mean feedforward input, due to the dominance of inhibition in our circuit. The anti-phase inhibition amplifies the modulation of the LGN input, yielding a stronger modulation of the net feedforward input. We previously showed that this combination of net inhibitory mean feedforward input, which is untuned for stimulus orientation, and the orientation tuning of the modulation of the feedforward input could account for the contrast-invariance of cortical orientation tuning (Troyer et al. 1998). Here we show that the same principles can explain suppression effects such as cross-orientation inhibition.

Finally, we introduce our terminology. We characterize a response to a periodic stimulus by the mean response, referred to as the "DC" of the response; and the amplitude of the response modulation at the temporal frequency of the stimulus, referred to as the "F1" of the response (for "first harmonic"). Note that the *peak* value of the response is roughly given by the sum DC+F1.

Cross-orientation inhibition

In the model circuit, the response to a base grating at the preferred orientation is suppressed by simultaneous presentation of a mask grating at the null orientation (the orientation perpendicular to the preferred). Two factors contribute to this suppression (Fig. 2). First, the preferred-orientation stimulus evokes synchronized modulation of the firing rates of the LGN inputs to a simple cell, but superposition of the null stimulus disrupts this synchronization. This reduces the F1 of the net LGN input to a simple cell. This is a small effect. Second, superposition of the null stimulus increases the mean spiking rate of the LGN inputs and thus the DC of the net LGN input to a simple cell. This is converted by the dominant antiphase inhibition in the circuit into an increased *negative* feedforward DC input to the simple cell. The combination of these two effects provides sufficient suppression to match many aspects of the suppression observed experimentally, including the illustrated cross-orientation inhibi-

tion. Note that the superposition of the second grating yields an increase in the peak input (DC+F1) of the LGN input, so that simple cell responses would be increased by addition of the second grating if only LGN input were considered.

Contrast dependence of cross-orientation inhibition

The amount of suppression is dependent on the mask contrast. The mean spike response to a base grating at the preferred orientation decreases with the contrast of a mask grating at a non-preferred orientation, both in a cell from Bonds (1989) (Fig. 3A) and in simulations (Fig. 3C). The curve of response vs. base contrast is shifted downward or rightward with increasing mask contrast, with little change in slope, both in an experimental cell (Fig. 3B) and in simulation (Fig. 3D). Both response vs. base contrast and suppression vs. mask contrast are stronger in the experimental cells illustrated than in simulation, but the degree of suppression for a given ratio of base to mask contrast is more comparable.

Place Fig. 3 about here

Recent experiments by Sengpiel et al. (1998) also addressed the effect of inhibition by a second grating on the contrast-response curve. They described the contrast-response function by a hyperbolic ratio function, $R = R_{max} \frac{c^n}{c_{50}^n + c^n}$, where R_{max} is the maximal response, c_{50} is the contrast that elicts a half-maximal response, and n is the power exponent. For a majority of cells, cross-orientation inhibition caused a rightward shift of the contrast-response function, that is, a change in c_{50} but not in R_{max} or n. This is seen in the normalized population response for 48 cells (both simple and complex cells), with and without cross-orientation inhibition, from Sengpiel et al. (1998) (Fig. 4A). The inhibition in the model is similarly described well by a rightward shift of the contrast response function (Fig. 4B).

Place Fig. 4 about here

Effect of cross-orientation inhibition on orientation tuning

We examined the effects of a mask stimulus at the null orientation on the response tuning for the orientation of a base grating (Fig. 5). The mask-induced suppression does not significantly alter the orientation tuning in the model: when the response tuning curve in the presence of the mask is scaled to have the same height as the response tuning curve in

Place Fig. 5 about here

the absence of the mask, the two curves become almost identical. Thus, the effect of the suppressing grating on the orientation tuning curve is predicted to be mainly divisive.

Dependence of suppression on mask orientation

The initial concept of cross-orientation inhibition has been broadened by several experiments to that of a more general, nonspecific suppression of the response to a preferred stimulus by superposition of a second stimulus (Bonds 1989; DeAngelis et al. 1992). Here we focus on the experiments of Bonds (1989), who used as stimuli superimposed pairs of sine gratings: a base grating at the preferred orientation and 2 Hz temporal frequency, and a mask grating at 3 Hz (or other temporal frequencies, considered later) and various orientations, as illustrated in figure 6. By decomposing the cell's response into two F1's, one at each temporal frequency, *Place* as well as a DC, the separate response to each grating was assessed: the F1 at 2 Hz represents *Fig. 6* response to the base grating, while the F1 at 3 Hz represents response to the mask grating. *about* This allowed Bonds (1989) to determine the effect of the orientation of the mask grating on *here* the suppression of the response to the base grating.

We have mimicked this procedure in our simulations (Fig. 7). In both experiments *Place* and simulations, the mask grating at non-preferred orientations suppressed the DC spike *Fig. 7* response, and the tuning of this DC response with mask orientation is similar to the exci- *about* tatory orientation tuning curve to a single grating (Fig. 7A,E). Similarly, both experiment *here* and simulation show tuning of the 3 Hz (mask) F1 component of the response with the mask orientation, again following the excitatory orientation-tuned response of the cell (Fig. 7B,F). The effects of the mask on the 2 Hz (base) F1 component of the response is more complex. Experimentally, Bonds (1989) found 11 of 14 cells showed suppression of the base component that was untuned for mask orientation (Fig. 7C), while 3 of 14 cells showed broad tuning for mask orientation (Fig. 7D). In simulations, we find that the 2 Hz component of the intra-cellular current shows suppression that is untuned for mask orientation (Fig. 7G), but the 2 Hz component of the spiking response shows broadly tuned suppression (Fig. 7H) akin to that seen in a minority of cells by Bonds. This tuning of the base spike response arises from

the addition of an untuned base current component and a tuned mask current component followed by a rectification to give the spike response.

Neither simulation results nor experimental results (Bonds 1989) depended on the relative spatial phase of the two gratings (not shown).

Temporal frequency dependence of the suppression

By keeping the base grating at 2 Hz and varying the temporal frequency of the mask grating, Bonds (1989) investigated the temporal characteristics of the cortical suppression. We first examine the dependence of suppression on mask contrast and temporal frequency (Figure 8). *Place* In the experiments, strong suppression is found for lower temporal frequencies (2 and 8 Hz), *Fig. 8* with a decrease in suppression at 16 Hz. In simulations, 8 and 16 Hz masks evoke greater *about* suppression than the 2 Hz mask, but the suppression does decrease between 8 and 16 Hz. *here*

We next examine the dependence of suppression on mask orientation and temporal frequency (Fig. 9). In experiments, the total (DC) response (Fig. 9A) shows suppression for *Place* masks at off-orientations that decreases with mask temporal frequency until, for a mask of *Fig. 9* 32 Hz, there is no inhibitory effect at all. When the mask is at the preferred orientation, the *about* lowest frequencies evoke no suppression. The mask response (Fig. 9B) is, not surprisingly, *here* tuned with mask orientation, and the excitation decreases with temporal frequency. The excitation actually cuts off earlier than the suppression: at 16 Hz there is little excitatory effect of the mask, while there is still a significant inhibition of the DC and base responses. The suppression of the base response (Fig. 9C) is untuned with mask orientation, and the inhibition is highest for low temporal frequencies and decreases with increasing temporal frequencies, as in Fig. 8. Our simulations (Fig. 9D-F) are similar except that, again, 3 Hz gratings evoke less suppression than higher temporal frequencies, and the base component of the spiking response shows orientation tuning (although again, the base component of the current response is untuned).

To understand the temporal-frequency tuning of suppression in our model, we examined the thalamocortical and full-circuit input to cells when adding a varying-temporal-frequency

mask grating at the orientation orthogonal to a preferred-orientation 2 Hz base grating (Fig. 10). We assume that the 2 Hz component of the response is roughly determined by the sum of the DC and the 2 Hz F1, representing the peak reached by the 2 Hz component when ignoring the other modulating components. The LGN input shows an increased DC component for 8 and 16 Hz masks relative to that for a 3 Hz mask, while the 2 Hz F1 component is slightly reduced. The cortical circuit amplifies both the DC and 2 Hz F1 component and reverses the DC component. The superposition of the 2 Hz F1 and DC then results in a peak input (F1+DC) that is lower for the higher-temporal-frequency masks.

Place Fig. 10 about here

According to the experiments, the peak input should increase for the higher mask temporal frequencies, suggesting flaws in our models either of the LGN DC and/or F1 or of inhibitory cell temporal tuning. However, after this work was completed, we learned of recent experiments (Bonds, manuscript submitted; Durand et al. 2001) suggesting that the temporal tuning of suppression actually is more like that predicted by the model; we return to this in the Discussion.

Spatial-frequency dependence of the suppression

Cortical cells show bandpass tuning for the spatial frequency of a grating. This bandpass tuning is expected to also influence suppressive effects. Figure 11 shows experiments (Bonds 1989) and simulations of cell responses with varying spatial frequency of the base and mask. The experiments show that tuning of the suppression for mask spatial-frequency is broader than the cell's excitatory spatial tuning bandwith. Furthermore the bandwith of the suppression increases with mask contrast. In our simulations, inhibitory cells and excitatory cells have similar spatial frequency tuning (they have identical Gabor-functions defining their receptive fields), and as a result, the suppression and excitatory tuning have similar bandwidth, suggesting a flaw in our model of inhibitory cell receptive fields. However, the simulations do show some broadening of the tuning of suppression with contrast.

Place Fig. 11 about here

Modulation changes with multiple gratings

When stimulating simple cells with a superposition of a high- and a low-temporal-frequency grating, the F1 of the spike response at the lower frequency is decreased relative to that evoked by the low temporal frequency grating alone. At the same time the F1 at the higher temporal frequency is enhanced compared to its value for a single grating. This was shown by Dean et al. (1982), who obtained results for 18 simple cells at 1.25 and 7.75 Hz temporal frequency using counter-phase gratings. When restricted to cells for which the ratio of the response amplitudes to each frequency presented alone was between 0.8-1.2, they found a mean relative response modulation (ratio of F1 when two gratings shown together to F1 of single grating alone) of 0.77 at 1.25 Hz and 1.23 at 7.75 Hz. Across all cells, they found mean changes in the same directions (Dean et al. 1982, fig 2). In our simulations we use drifting rather than counterphase gratings so that the modeled responses will not depend on the absolute spatial phase of the cells, allowing cells of all absolute spatial phases to be studied simultaneously. Using 2 Hz and 8 Hz gratings with contrasts chosen to yield similar responses to each grating alone, we find a relative response modulation of 0.60 at 2 Hz and 1.09 at 8 Hz (Fig. 12A). Thus the strength of suppression and the direction, though not the strength, of enhancement are roughly replicated by our model.

Place Fig. 12 about here

This can be understood from a simple toy model as illustrated in figure 13. We assume that each grating alone evokes both a sinusoidal oscillation and a negative DC; the superposition of the gratings is modeled simply by adding their oscillations and DC's. The resulting waveform is rectified at a threshold and the F1 of the rectified response is computed. The model has four parameters: the threshold θ, expressed as a percentage of the amplitude of the 2 Hz oscillation; the sizes β_{2Hz} and β_{8Hz} of the negative DC's, relative to the amplitude of their respective sinusoidal oscillations; and the amplitude α of the 8 Hz oscillation relative to that of the 2 Hz. This toy model is similar to one considered by Dean et al. (1982) (their Fig. 1), except that their model did not include the negative DC's, *i.e.* in their model $\beta_{2Hz} = \beta_{8Hz} = 0$. They found that their toy model could *not* explain their results, and indeed we shall find that the negative DC's are critical to our explanation of the results.

Place Fig. 13 about here

In simulations of the full model, we find $\beta_{8Hz} = 0.4232$ and $\beta_{2Hz} = 0.0187$ (see Methods), and so accordingly in our toy model we initially consider $\beta_{8Hz} = 0.4$ and $\beta_{2Hz} = 0$. We first choose α so that, after rectification, the 2 Hz grating alone and the 8Hz grating alone evoke similar F1's. Examining results across values of the threshold (Fig. 12B), we find that for thresholds less than 0.3, this simple model replicates the basic result of the experiment: the low-frequency F1 decreases, and the high-frequency F1 increases, when the two sinusoids are added together. Next we fix the threshold at 0.2, and find that across all values of α, the same basic result is obtained (Fig. 12C).

Finally we examine the dependence of the results on the two β's, that is, on the size of the negative DC offset evoked by the two gratings, for fixed values of the threshold and relative amplitude (Fig. 12D-F). Small or negative values of β_{2Hz} and high (positive) values of β_{8Hz} yield changes in the experimentally observed directions: for these values of the β's, the 2 Hz F1 is decreased (negative "2 Hz F1 difference" in Fig. 12D) and the 8 Hz F1 is increased (positive "8 Hz F1 difference" in Fig. 12E) when the two sinusoids are added together. To better ascertain the parameter range over which each F1 is changed in the appropriate direction, we multiply the negative region of the 2 Hz change (positive values set to zero) with the positive region of the 8 Hz change (negative values set to zero), and take the absolute value of the result; contours of this result are shown in Fig. 12F. The contours increase with approach toward the lower right corner, that is, as β_{8Hz} increases and β_{2Hz} decreases. (The parameter region that gives the desired result increases when the threshold is decreased or the relative 8 Hz amplitude is increased, not shown.) Thus, we conclude that it is critical to the observed result that an 8 Hz grating evoke a larger negative DC, relative to its F1, than a 2 Hz grating. (This occurs in our full model due to the slow time course of NMDA conductances, which lowers the F1/DC ratio for higher frequencies; see Discussion.)

Reid et al. (1992) studied the superposition of 8 counter-phase gratings of increasing temporal frequencies. As in the two-grating paradigm of Dean et al. (1982), they found that the modulation of low temporal frequencies is depressed and that of high temporal frequencies is enhanced in the superposition relative to that of the constituent sinusoids.

We again use drifting gratings, and examine the F1 responses of our full model at each of 8 stimulus frequencies, in response to each of the 8 individual sinusoids alone and in response to their superposition (Fig. 14). As in the experiments, the F1's at lower temporal frequencies are decreased and those at higher temporal frequencies are increased in the superposition, relative to the single gratings. We find a cross-over point from depression to elevation between 4 and 5 Hz, slightly lower than, but comparable to, the experimentally found cross-over of 6 to 8 Hz (Reid et al. 1992). *Place Fig. 14 about here*

Discussion

We have shown that a simple correlation-based local circuit can account for many aspects of cortical responses to multiple gratings, including many features of the dependence of suppression or enhancement on a mask grating's orientation, spatial frequency, and temporal frequency, as well as the effects of mixing sinusoids of multiple temporal frequencies. The model predicts that superposition of a mask grating orthogonal to the preferred orientation should cause little change in orientation tuning to a base grating, having a primarily divisive effect on the orientation tuning curve; to our knowledge this has not been tested. It also predicts that high-temporal-frequency (*e.g.*, 8 Hz) gratings should evoke a stronger negative DC for a given F1 amplitude than low-temporal-frequency gratings (*e.g.*, 2 Hz), and that this larger DC at higher frequencies is critical to explaining the effects of mixing sinusoids of multiple temporal frequencies (see Figs. 12-13).

The circuit model that attains these results (Troyer et al. 1998) rests on three basic assumptions. First, both the excitation and inhibition received by a simple cell come from other cells of similar preferred orientation (Anderson et al. 2000a; Ferster 1986), but the excitation and inhibition are evoked by opposite polarities (light or dark) at any given point in the receptive field (Ferster 1988; Hirsch et al. 1998). Second, the inhibition dominates, rather than balances, excitation in the cortical circuitry. This dominance is supported by experiments showing that even slight movement of a spot stimulus across a subregion bound-

ary is sufficient to cause the spot to evoke net inhibition (Hirsch et al. 1998), as well as by experiments showing that nonspecific shock of the LGN evokes massive inhibition in cortex (Ferster and Jagadeesh 1992). It is this excess inhibition that primarily accounts for the cross-orientation (or 'second grating') suppression. Third, the inhibitory neurons in the model respond to all orientations, although they are orientation-tuned: their tuning curves resemble a tuned hill atop an untuned platform (Troyer et al. 1998). Recent results find two types of inhibitory interneurons in cat V1 layer 4: simple cells that show good orientation tuning, and complex cells that are untuned for orientation (Hirsch et al. 2000). It is possible that these separately embody the two components of inhibitory tuning – the tuned hill and the untuned platform – that we are attributing to a single class of cells.

The model points to two basic sources of second-grating suppression. In the model, the degree of activation of a cell is determined by the size of the temporal modulation or F1 of a simple cell's LGN input, relative to the size of the mean level or DC of this input: the F1 leads to activation, while the DC is transformed by the dominant antiphase inhibition into a net suppression. Thus, the second grating may suppress responses in either of two ways: (1) by lowering the F1 of a simple cell's LGN input, *e.g.* due to desynchronization of the rate modulations of the individual LGN inputs to a cell; and/or (2) by raising the DC level of the cell's LGN inputs inputs, evoking net inhibition. In our hands the change in F1 is a small effect, but our model of LGN response is currently very simple: we assume LGN spike-rate responses to the two gratings add linearly, up to rectification at zero firing rate. Furthermore, short-term synaptic depression of geniculocortical synapses, which we have not considered here, can lower the base F1 of the LGN input in the presence of a mask, as discussed below. Thus, it is possible that in reality this could be a larger effect. In our hands, the increase in DC and the resulting inhibition is the primary source of suppression. However, synaptic depression can suppress this DC at lower temporal frequencies, discussed below. Furthermore, Bonds (1989) shows results from a few LGN cells suggesting that the response to two gratings is sublinear relative to the response to each grating alone, with little change in DC. It will be important to better characterize both the degree to which the DC

of the LGN input grows, and the manner in which the F1 of the total LGN input received by a simple cell alters, when a mask grating is added to the base grating.

A paradigm we have not addressed here is that, when a cell is tonically excited by a one-dimensional noise stimulus at the preferred orientation $\mathcal{P}$, a counterphase grating at the null orientation $\mathcal{N} = \mathcal{P} + 90^o$ induces frequency-doubled suppression (Morrone et al. 1982). This has led to the suggestion that suppression may be mediated by complex cells, or a pool of simple cells of many preferred phases, preferring the null orientation $\mathcal{N}$ (Bonds 1989; DeAngelis et al. 1992; Morrone et al. 1982). We have not attempted to simulate either noise or counterphase stimuli, but it is worth noting that this frequency doubling is expected in our model. Each of the two opposite phases of the counterphase grating at angle $\mathcal{N}$ will excite roughly half the LGN inputs to an inhibitory cell preferring angle $\mathcal{P}$, and inhibit the other half. The LGN inputs that are excited can raise their firing rates much more than those that are inhibited can lower their firing rates (because lowering is bounded at zero firing rate), so the net result is that each phase produces a pulse of increased LGN input causing the inhibitory cell to respond. Because inhibition dominates excitation, the feedforward inhibition evoked by the pulses of LGN input will exceed the feedforward excitation they evoke, and a frequency-doubled suppression of excitatory cell responses will result.

A Role for Synaptic Depression?

Since this work was completed, Carandini and colleagues have reported that cross-orientation suppression has temporal frequency tuning like that of LGN inputs (Durand et al. 2001), and is not affected by a period of adaptation to a cross-oriented stimulus (Freeman et al. 2001). Based on these results and previous findings that cross-orientation suppression is largely, though not entirely, monocular (Walker et al. 1998), they proposed that cross-orientation suppression does not stem from the firing of cortical cells but rather from short-term synaptic depression of geniculocortical synapses (Durand et al. 2001; Freeman et al. 2001).

We had examined the effects of including such depression in our model (modeled as in the "pulse" parameters of Kayser et al. (2001)), and found that it does not substantially

alter our results as to the contrast, temporal, spatial, and mask orientation tuning of mask suppression or the effects of combining two or more gratings of different temporal frequencies (unpublished data). However, we have now reexamined this and found that inclusion of depression does alter the *mechanism* of suppression: in the presence of synaptic depression, the base F1 component of the LGN input is substantially decreased by the presence of a mask. Synaptic depression also substantially lowers the DC of the LGN input to low temporal frequencies (*e.g.* 3 Hz), but not to higher temporal frequencies (*e.g.* 8 Hz) (Krukowski 2000). The result is that, when the base grating is combined with a low-temporal-frequency mask, the DC is near zero so the inhibition plays little role, and most of the mask suppression is due to the lowering of the base F1 by synaptic depression. However, when the base grating is combined with a high-temporal-frequency mask, the DC is substantial, and so the combination of the suppression of the base F1 by depression and the reversal of the DC by antiphase inhibition is needed to explain mask suppression.

A strong argument for a role of inhibition is that, when a cell's firing rate is elevated by iontophoresis of an amino acid, a cross-oriented stimulus reduces responses (Ramoa et al. 1986). This has no obvious explanation in terms of synaptic depression of geniculocortical synapses, but is easily explained by inhibition.

An important distinction between feedforward inhibition and synaptic depression as mechanisms mediating mask suppression is that the former should have essentially instantaneous effects, while the latter should take some time to develop. While feedforward inhibition involves an extra synapse relative to LGN excitation, inhibitory cells have lower thresholds and are the first to spike in response to LGN excitation, so that feedforward inhibition arrives within a few milliseconds of the arrival of LGN excitation (*e.g.*, Ferster and Jagadeesh 1992). However, it is possible that the effects of geniculocortical depression become substantial over the time in which a cortical cell integrates its inputs before spiking, so that depression effects also would appear to be instantaneous in extracellular recording. Both theoretical studies of the time course of mask suppression expected from synaptic depression, and experimental characterization of the time course of mask suppression, will be of great interest.

The finding that adaptation to a cross-oriented stimulus does not alter the strength of cross-orientation suppression (Freeman et al. 2001) suggests that the layer 4 inhibitory neurons in our model should not show contrast adaptation, or at least not in response to a cross-oriented stimulus. To our knowledge, the presence of contrast adaptation in inhibitory neurons has never been tested; it would be very interesting to test this prediction. However, it will also be important to determine whether adaptation affects the strength of cross-orientation suppression induced by higher-temporal-frequency masks, since it is these masks that should elicit the strongest inhibition in the presence of synaptic depression.

Failures of the Model

While the model captures a wide array of results, it also fails in two or three notable ways. One apparent failure is that suppression in the experiments of Bonds (1989) is tuned to lower mask temporal frequencies than in the model. However, since this work was completed, we learned of new experiments by Bonds (manuscript submitted) and by Durand et al. (2001), both reporting the high-temporal frequency cutoff of suppression to be in the range of 16-20 Hz, similar to our findings (also seen in Morrone et al. 1982). Thus, what appeared to be a flaw in the model may instead turn out to be an accurate prediction. This high cutoff arises in our model because the temporal frequency tuning of the inhibitory neurons follows that of the LGN input.

This leaves two remaining problems. First, suppression in experiments is more broadly tuned for mask spatial frequency than in the model (also seen in DeAngelis et al. 1992; Morrone et al. 1982). Second, the majority of cells studied experimentally showed no tuning for mask orientation of the base component of response, whereas in the model there is such tuning.

It will be important to determine these properties for cells in layer 4, the layer that we are modeling. If they hold in layer 4, it would suggest that two types of improvements are needed in the model. First, we may need to consider inhibitory cell models that have broader tuning or a broader variety of tunings for spatial frequency. The model presently assumes

that inhibitory neurons have the same multiple-subregion Gabor receptive field shape as excitatory neurons. A source for broader tuning was suggested by Toyama et al. (1981), who reported that inhibition onto layer 4 simple cells with multiple ON/OFF subregions may come primarily from inhibitory layer 4 simple cells with single-subregion receptive fields. Second, the tuning for mask orientation of the base response component is a threshold effect: the addition of the untuned base component and the tuned mask component of current, followed by rectification by the spike threshold, yields a tuned base component of the spiking response. Sufficient voltage noise can smooth away many threshold effects, and voltage noise appears comparable in size to stimulus-induced voltage modulations (Anderson et al. 2000b; Arieli et al. 1996; Paré et al. 1998; Tsodyks et al. 1999), suggesting that exploration of higher-noise regimes might remove the threshold-induced tuning.

Suppression and contrast

The modeled experiments (Bonds 1989) and other reported results (DeAngelis et al. 1992; Morrone et al. 1982) display a wide range of suppression strengths. There does not seem to be any systematic connection between suppression strengths and cell response properties or cortical position, except that simple cells are reported to show stronger suppression than complex cells (Morrone et al. 1982). Our cells often required higher contrasts of both base and mask than in experiments to see a given level of response or suppression, although the relative contrast of base vs. mask for a given strength of suppression was often similar. Strength of overall responses can be varied without otherwise altering circuit behavior by varying circuit parameters in a coordinated way (Troyer et al. 1998); thus the most robust issue seems to be the relative contrast required to achieve a given degree of suppression, rather than the absolute contrasts. Furthermore, the relative contrast necessary to achieve a given degree of suppression is also parameter-dependent in the model, depending in particular on the overall strength of inhibition. We have used the model as tuned for other experiments and not altered it to fit these data. The strength of response and suppression also depends on our LGN model, which as we have noted is very simple and needs better characterization.

Modulation changes with multiple gratings

Our model reproduces the qualitative effects of modulation changes for pairs of gratings found by Dean et al. (1982). The simple toy model we considered for this paradigm accords well with this, showing the experimentally observed effect across parameters provided that the 8 Hz grating evokes a stronger negative DC for a given F1 than the 2 Hz grating. This stronger negative DC at higher frequencies arises in our full model due to the presence of NMDA conductances in excitatory synapses; the slow time course of these conductances suppresses the high-frequency F1 without altering the DC, resulting in a higher DC/F1 ratio (*e.g.* Krukowski 2000, Fig. 2.28). Frequency-dependent synaptic depression in geniculocortical synapses can play a similar role, by differentially lowering the DC/F1 ratio of low-frequency stimuli (Krukowski 2000).

Our result for the extension to 8 sine gratings, Fig. 14, agrees with the experimentally observed response amplitudes of cortical cells (Reid et al. 1992). Reid et al. (1992) further reported a decrease in integration time of the response to the 8 sine gratings, as determined by the slope of the curve of response phase vs. temporal frequency; the cells appear to respond faster. They argue that these effects must be of cortical origin. We do not see a change in integration time in our simulations (data not shown). We have found (Kayser et al. 2001) that increases in contrast can advance the phase of response, due to a number of mechanisms including synaptic depression, spike-rate adaptation, and increases in conductance. We can speculate that such effects might contribute to the experimental result, since there is greater overall contrast in the 8-grating stimulus than in the single-grating stimuli, although we do not see an effect in the present simulations.

Conclusion

Correlation-based or "push-pull" connectivity with dominant inhibition has been shown to account for many experimentally observed properties of layer 4 of cat primary visual cortex. These include contrast-invariant orientation tuning (Troyer et al. 1998), cortical temporal frequency tuning that cuts off at much lower frequency than LGN tuning (Krukowski and

Miller 2001), and various contrast-dependent nonlinearities (Kayser et al. 2001). The present results suggest that this framework can also explain a range of suppression and enhancement effects observed with multiple-grating stimuli. However, they also suggest some modifications may be needed: in particular, our inhibitory cell models may need to be modified to explain the broader spatial-frequency tuning of suppression observed experimentally, and we may need to explore regimes with more realistic levels of voltage noise.

Acknowledgements

References

Alonso, J. M., Usrey, W. M, and Reid, R. C. Precisely correlated firing in cells of the lateral geniculate nucleus. *Nature* 383:815–819, 1996.

Anderson, J. S., Carandini, M., and Ferster, D. Orientation tuning of input conductance, excitation, and inhibition in cat primary visual cortex. *J. Neurophysiol.* 84:909–926, 2000a.

Anderson, J. S., Lampl, I., Gillespie, D., and Ferster, D. The contribution of noise to contrast invariance of orientation tuning in cat visual cortex. *Science* 290:1968–1972, 2000b.

Arieli, A., Sterkin, A., Grinvald, A., and Aertsen, A. Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* 273:1868–1871, 1996.

Bonds, A. B. Role of inhibition in the specification of orientation selectivity of cells in the cat striate cortex. *Vis. Neurosci.* 2:41–55, 1989.

Carandini, M. and Heeger, D. J. Summation and division by neurons in visual cortex. *Science* 264:1333–1336, 1994.

Carandini, M., Heeger, D. J., and Movshon, J. A. Linearity and normalization in simple cells of the macaque primary visual cortex. *J. Neurosci.* 17:8621–8644, 1997.

Carandini, M., Heeger, D. J., and Movshon, J. A. Linearity and gain control in V1 simple cells. In: *Cerebral Cortex, Vol. 13: Models of cortical circuits*, edited by P.S. Ulinski, E. G. Jones, and A. Peters. Kluwer Academic/Plenum New York, 1999, pp. 401–443.

Carmignoto, G. and Vicini, S. Activity-dependent decrease in NMDA receptor responses during development of the visual cortex. *Science* 258,:1007–1011, 1992.

Chung, S. and Ferster, D. Strength and orientation tuning of the thalamic input to simple cells revealed by electrically evoked cortical suppression. *Neuron* 20:1177–89, 1998.

Crair, M. C. and Malenka, R. C. A critical period for long-term potentiation at thalamocortical synapses. *Nature* 375:325–328, 1995.

Dean, A. F., Tolhurst, D. J., and S., Walker N. Non-linear temporal summation by simple cells in cat striate cortex demonstrated by failure of superposition. *Exp. Brain Res.* 45:456–458, 1982.

DeAngelis, G. C., Robson, J. G., Ohzawa, I., and Freeman, R. D. Organization of suppression in receptive fields of neurons in cat visual cortex. *J. Neurophysiol.* 68:144–163, 1992.

Durand, S., Freeman, T. C. B., Kiper, D. C., and Carandini, M. Probing the sources of cross-orientation suppression in cat visual cortex using fast stimuli. *Perception (suppl.)* (abstract), in press, 2001.

Ferster, D. Orientation selectivity of synaptic potentials in neurons of cat primary visual cortex. *J. Neurosci.* 6:1284–1301, 1986.

Ferster, D. Spatially opponent excitation and inhibition in simple cells of the cat visual cortex. *J. Neurosci.* 8:1172–1180, 1988.

Ferster, D., Chung, S., and Wheat, H. Orientation selectivity of thalamic input to simple cells of cat visual cortex. *Nature* 380:249–252, 1996.

Ferster, D. and Jagadeesh, B. EPSP-IPSP interactions in cat visual cortex studied with in vivo whole-cell patch recording. *J. Neurosci.* 12(4):1262–1274, 1992.

Ferster, D. and Miller, K. D. Neural mechanisms of orientation selectivity in the visual cortex. *Ann. Rev. Neurosci.* 23:441–471, 2000.

Freeman, T. C. B., Durand, S., Kiper, D. C., and Carandini, M. Probing the source of cross-orientation suppression in primary visual cortex. *Inv. Opth. Vis. Sci.* 42:S726, 2001.

Heeger, D. J. Normalization of cell responses in cat striate cortex. *Vis. Neurosci.* 9:181–198, 1992.

Heeger, D. J., Simoncelli, E. P., and Movshon, J.A. Computational models of cortical visual processing. *Proc. Natl. Acad. Sci. USA* 93:623–627, 1996.

Hirsch, J. A., Alonso, J.-M., Pillai, C., and Pierre, C. Simple and complex inhibitory cells in layer 4 of cat visual cortex. *Soc. Neurosci. Abstr.* 26:1083, 2000.

Hirsch, J. A., Alonso, J.-M., Reid, R. C., and Martinez, L.M. Synaptic integration in striate cortical simple cells. *J. Neurosci.* 18:9517–9528, 1998.

Hubel, D. H. and Wiesel, T. N. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.* 160:106–154, 1962.

Jahr, C. E. and Stevens, C. F. Voltage dependence of NMDA-activated macroscopic con-

ductances predicted by single-channel kinetics. *J. Neurosci.* 10:3178–3182, 1990.

Kayser, A. S., Priebe, N. J., and Miller, K. D. Contrast-dependent nonlinearities arise locally in a model of contrast-invariant orientation tuning. *J. Neurophysiol.* 85:2130–2149, 2001.

Krukowski, A. E., 2000. *A model of cat primary visual cortex and its thalamic input.* PhD thesis University of California, San Francisco.

Krukowski, A. E. and Miller, K. D. Thalamocortical NMDA conductances and intracortical inhibition can explain cortical temporal tuning. *Nature Neuroscience* 4:424–430, 2001.

Lester, R. A., Clements, J. D., Westbrook, G. L., and Jahr, C. E. Channel kinetics determine the time course of NMDA receptor-mediated synaptic currents. *Nature* 346:565–567, 1990.

Morrone, M. C., Burr, D. C., and Maffei, L. Functional implications of cross-orientation inhibition of cortical visual cells. I. Neurophysiological evidence. *Proc. R. Soc. Lond. B* 216:335–354, 1982.

Paré, D., Shink, E., Gaudreau, H., Destexhe, A., and Lang, E. J. Impact of spontaneous synaptic activity on the resting properties of cat neocortical pyramidal neurons in vivo. *J. Neurophysiol.* 79:1450–1460, 1998.

Ramoa, A. S., Shadlen, M., Skottun, B. C., and Freeman, R. D. A comparison of inhibition in orientation and spatial frequency selectivity of cat visual cortex. *Nature* 321:237–239, 1986.

Reid, R. C. and Alonso, J. M. Specificity of monosynaptic connections from thalamus to visual cortex. *Nature* 378:281–284, 1995.

Reid, R. C., Victor, J. D., and Shapley, R. M. Broadband temporal stimuli decrease the integration time of neurons in cat striate cortex. *Vis. Neurosci.* 9:39–45, 1992.

Sclar, G. Expression of "retinal" contrast gain control by neurons of the cat's lateral geniculate nucleus. *Exp. Brain Res.* 66:589–596, 1987.

Sclar, G. and Freeman, R. D. Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Exp. Brain Res.* 46:457–461, 1982.

Sengpiel, F., Baddeley, R. J., Freeman, T. C. B., Harrad, R., and C., Blakemore. Different mechanisms underlie three inhibitory phenomena in cat area 17. *Vision Res.* 38:2067–2080,

1998.

Skottun, B. C., Bradley, A., Sclar, G., Ohzawa, I., and Freeman, R. D. The effects of contrast on visual orientation and spatial frequency discrimination: A comparison of single cells and behavior. *J. Neurophysiol.* 57:773–786, 1987.

Tanaka, K. Cross-correlation analysis of geniculostriate neuronal relationships in cats. *J. Neurophysiol.* 49:1303–1318, 1983.

Toyama, K., Kimura, M., and Tanaka, K. Organization of cat visual cortex as investigated by cross-correlation technique. *J. Neurophysiol.* 46:202–214, 1981.

Troyer, T. W., Krukowski, A., Priebe, N. J., and Miller, K. D. Contrast-invariant orientation tuning in cat visual cortex: Feedforward tuning and correlation-based intracortical connectivity. *J. Neurosci.* 18:5908–5927, 1998.

Tsodyks, M. V., Kenet, T., Grinvald, A., and Arieli, A. Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science* 286:1943–1946, 1999.

Walker, G. A., Ohzawa, I., and Freeman, R. D. Binocular cross-orientation suppression in the cat's striate cortex. *J. Neurophysiol.* 79:227–239, 1998.

Figure Legends

Figure 1

A: Cartoon figure of the correlation-based local cortical circuitry used in our model. Four cell pools – two excitatory (top) and two inhibitory (bottom), each with two opposite preferred spatial phases (left vs. right) – are depicted. The four cell pools are at the same retinotopic position, but have been separated for visibility. Light gray represents ON subfields and dark grey OFF subfields. Cells receive input from LGN and from other cortical cells preferring similar orientations. Connections are assigned probabilistically. ON-center LGN inputs with centers overlying a cell's ON-subregions, and OFF inputs with centers overlying OFF subregions, are likely to connect to the cell. The cells receive excitatory intracortical connections with highest probability from other excitatory cells they are most correlated with (same absolute placement of ON and OFF subfields). Inhibitory connections are received with highest probability from inhibitory interneurons they are most anti-correlated with (opposite absolute placement of ON and OFF subfields). Unlike the cartoon, the actual model includes cells of all orientations and spatial phases and spanning a range of retinotopic positions. The probabilistic sampling leads cells to receive input from other cells differing in preferred orientation by up to about 30^o, and differing in phase from identical (excitatory connections) or opposite (inhibitory connections) by up to about 60^o; the cartoon illustrates the most probable connections. The feedforward inhibition (LGN→inhibitory cell→ excitatory cell) is stronger than the feedforward excitation (LGN→excitatory cell).

B: Membrane potential (top) and synaptic currents (bottom) of a single cell as a function of time when stimulated by a drifting grating at its preferred orientation. The excitatory (black) and inhibitory (gray) synaptic currents are modulated out of phase with one another, resulting in an oscillatory membrane potential. The cell spikes during the excitatory peak of the oscillations. C: Membrane potential (top) and synaptic currents (bottom) of a single cell when stimulated by a drifting grating orthogonal to its preferred orientation. There is little modulation of the synaptic current or membrane potential in time. Since the feedforward

inhibition is stronger than the feedforward excitation, the membrane potential does not reach threshold.

Figure 2

F1, DC and peak (F1+DC) input current to a model simple cell (mean over cells with similar orientation preference, 4 degree bins) from LGN alone (A) and LGN + cortex (B) in response to stimulus grating at preferred orientation, 40% contrast alone (base); or to superposition of preferred (base) and orthogonal (mask) orientation gratings, each at 40% contrast.

Figure 3

Experiments reproduced from Bonds (1989) (A, B – each shows data for a single simple cell), and corresponding simulation results (C, D – mean over cells with similar orientation preference; error bars in C show standard deviation, error bars omitted in D for visual clarity). A, C: Suppression of mean spiking response to a base grating of optimal orientation (expt. 28% contrast, sim. 40% contrast) with varying contrast of mask at non-preferred orientation ("inhibiting" orientation in experiments; orthogonal to preferred in simulations). Dashed line is response with base alone. B, D: Response versus log contrast of base shows downward shift that increases with mask contrast, with little change in slope. Note, overall spiking rate differences are not meaningful: experimentally there is high variability between cells (compare A,B); in simulations, overall spiking rate can be modified by parameter changes without otherwise altering network behavior (Troyer et al., 1998).

Figure 4

Experiments reproduced from Sengpiel et al. (1998) (A – mean response of pool of 48 cells, simple and complex from all layers, each cell's response expressed relative to its maximal response), and corresponding simulation results (B – mean spike response, ± 1 stdev, over cells in simulation with similar orientation preference). Response R vs. base contrast c fitted

to hyperbolic rate functions $R = R_{\max}c^n/(c^n + c_{50}^n)$, shown as dashed lines (control responses, base grating alone) and dotted lines (base plus mask gratings). Experiments, control: $R_{\max}$ = 1.072, c_{50} = 0.118, n = 1.46 (χ^2 = 0.097, 3 degrees of freedom (d.f.)); plus cross-oriented mask: c_{50} = 0.308 with $R_{\max}$ and n held fixed (χ^2 = 2.49, 5 d.f.). Simulations, control: $R_{\max}$ = 42.8013, c_{50} = 0.8564, n = 1.126 (χ^2 = 0.1131, 2 d.f.); plus cross-oriented mask (40 % contrast): c_{50} = 1.5707 with $R_{\max}$ and n held fixed (χ^2 = 3.6672, 4 d.f.). For both experiments and simulations, cross-orientation inhibition is well described by a rightward shift of the contrast response function.

Figure 5

Model orientation tuning curves, with and without mask stimulation. Solid line is orientation tuning curve when stimulated with a single base grating alone (mean ± 1 SD). Dashed line is the tuning (mean) for base grating orientation while simultaneously stimulating with a mask grating at the orientation orthogonal to the cell's preferred orientation. Dotted line is the same as the dashed, except scaled so that the peak response matches that of the tuning curve for the base grating alone. The response reduction due to the mask grating is mainly multiplicative.

Figure 6

Experimental procedure of Bonds (1989). Bonds created compound gratings by adding two simple gratings: a base grating always at the preferred orientation with a temporal frequency of 2 Hz; and a mask grating at various orientations with a temporal frequency of 3 Hz (in later experiments the mask has other temporal frequencies). He then examined the mean (DC) response, the base response defined as the F1 response at 2 Hz, and the mask response defined as the F1 response at 3 Hz.

Figure 7

Experiments reproduced from Bonds (1989) (A, B, C: data for a single simple cell, dashed line is base, 10% contrast at preferred orientation, dotted line is with addition of a 10% mask grating, and solid line is with addition of a 20% mask grating at various orientations; D: data from another single simple cell, dashed line is base, 7% contrast at preferred orientation, dash-dotted line is with addition of a 14% mask grating), and corresponding simulation results (E, F, G, H: each shows mean response over cells in simulation with similar orientation preference; dashed line is base, 40% contrast at preferred orientation, dotted line is with addition of a 20% mask grating, and solid line is with addition of a 40% mask grating at various orientations). A, E: Mean spike response vs. mask orientation. B, F: 3 Hz F1 (mask component) of the cell response is similar for experiment and simulations. C: 2 Hz F1 (base component) of the cell from A and B is untuned with the mask orientation (type I, 11/14 simple cells studied). D: 2 Hz F1 component of another simple cell shows some tuning with the mask orientation (type II, 3/14 simple cells). G: 2 Hz F1 component of the synaptic current in simulations is untuned with mask orientation, similar to type I cells. H: 2 Hz F1 component of spike response in simulations shows tuning with the mask orientation, similar to type II cells.

Figure 8

Experiments reproduced from Bonds (1989) (A: data for a single simple cell), and corresponding simulation results (B: mean response over cells in simulation with similar orientation preference). Plots show mean spike response vs. mask contrast for a combination of a 2 Hz base grating at preferred orientation and a mask grating oriented at an inhibitory orientation (experiments), which we have taken to be orthogonal to the preferred (simulations). Dashed lines, response to base grating alone; solid lines, response with 2 Hz mask; dash-dotted lines, response with 8 Hz mask; dotted lines, response with 16 Hz mask. The inhibitory effect of the mask grating decreases with the temporal frequency of the mask in experiments, whereas in the simulations the inhibitory effect of the mask grating is tuned

towards higher temporal frequencies, thus peaking at 8 to 16 Hz, but decreasing with further increase in temporal frequency (not shown).

Figure 9

Dependence of suppression on orientation and temporal frequency of mask grating, added to base grating at preferred orientation and 2 Hz. Experiments reproduced from Bonds (1989) (A, B, C: data for a single simple cell), and corresponding simulation results (D, E, F: mean response over cells with similar orientation preference). "Control" is response to base grating alone. A, D: Mean (DC) spike response. B, E: F1 of the spike response at the mask temporal frequency. C, F: F1 of spike response at the base temporal frequency. In general, inhibition peaks at higher temporal frequencies in the simulations than in the experimental cell. As earlier, the simulated base response is tuned with the mask orientation, while the experimental base response is not. Simulated base current response is untuned (not shown).

Figure 10

Base (2 Hz) F1, DC, and peak (2 Hz F1+DC) input current to a model simple cell from LGN alone (A) and LGN + cortex (B) in response to stimulus grating at preferred orientation, 40% contrast alone (base); or to superposition of preferred (base) and orthogonal orientation (mask) gratings of various temporal frequencies, each at 40% contrast. The DC input from LGN alone increases with higher frequencies up to 16 Hz, following the tuning of our model LGN cells (Sclar 1987). This results in an increased (inhibitory) DC for the full network which, added to a slightly smaller base modulation (F1), results in a significantly decreased peak input to the cell, thus explaining the temporal tuning of the inhibition of the base component of response in the simulations.

Figure 11

Experiments reproduced from Bonds (1989) (A: data for a single simple cell), and corresponding simulation results (B: each shows mean response over cells with similar orientation preference). Response vs. spatial frequency of base and mask; base is at cell's preferred orientation, mask is at an inhibitory orientation (experiment) or orthogonal to the preferred (simulation). Dashed lines with circles show response to base alone, vs. base spatial frequency; horizontal dashed lines show response to base alone at its preferred spatial frequency. Inhibitory curves show the response, vs. mask spatial frequency, when adding a mask at high (solid) and low (dash-dotted) contrast to the base at its preferred spatial frequency. The inhibition has broader spatial frequency tuning for the experiments than in the simulations.

Figure 12

A: F1 response at 2 and 8 Hz for a single drifting grating at the preferred orientation, at a temporal frequency of 2 Hz (20% contrast) and 8 Hz (40% contrast) respectively (white bars), compared with the 2 and 8 Hz F1 response to a superposition of the two gratings (both at the preferred orientation) (grey bars). The contrasts were chosen to give a similar F1 of the spiking response to the single gratings. The 2 Hz F1 response decreases when a higher temporal frequency grating is superposed with it, while the 8 Hz F1 response increases (slightly) when a lower temporal frequency grating is superposed with it. B-F: Results from a toy model of the effect in A; the toy model is illustrated in Fig. 13 and further described in Results. B-C: Solid: 2 Hz F1 response to 2 Hz grating alone; dash-dotted: 8 Hz F1 response to 8 Hz grating alone; dashed: 2 Hz F1 response to double grating; dotted: 8 Hz F1 response to double grating. Negative DC's are matched to data from the full model; $\beta_{2Hz} = 0$, $\beta_{8Hz} = 0.4$. B: $\alpha = 1.95$. For thresholds less than around 0.3, the model displays the biologically observed behaviour (Dean et al. 1982): the 2 Hz F1 for the double grating is less than for the single grating, while the 8 Hz F1 for the double grating is higher than for the single grating. C: Threshold is fixed at 0.2. For all values of α, the F1 responses are

modified in the double gratings in the directions observed biologically. For higher values of α, the double grating F1 responses approach their respective single grating F1 responses. D-F: Dependence of results on the negative DC's, β_{2Hz} and β_{8Hz}, for a fixed threshold ($\theta = 0.2$) and relative amplitude ($\alpha = 2$). D: Contour plot of the relative change in the 2 Hz F1 (F1 given both sinusoids minus F1 given 2 Hz sinusoid alone, divided by F1 for 2 Hz sinusoid alone). Solid lines are positive values and dashed lines are negative values, each line indicates an increment of 0.05 in relative values. The relative change is negative, as in the experiments, for low values of β_{2Hz} and high values of β_{8Hz}. E: Relative change in the 8 Hz F1, conventions as in D; this relative change is positive, as in the experiments, for low values of β_{2Hz} and high values of β_{8Hz}. F: Negative values in D are multiplied by positive values in E, and the absolute value of the result is shown as a contour plot, with values increasing toward lower right corner as indicated by lighter grays. The toy model behaves as in the experiments for low values of β_{2Hz} and high values of β_{8Hz}.

Figure 13

Toy model of the input and response to single gratings with temporal frequencies of 2 Hz (top) and 8 Hz (middle) and to the superposition of the two gratings (bottom). In each panel, the solid oscillating curves show the net inputs evoked by the gratings as a function of time. The input to the single gratings has two components, a component that modulates sinusoidally in time and a negative constant (mean or DC) component. The amplitude of the 8 Hz sinusoid, relative to that of the 2 Hz sinusoid, is α. The negative mean, shown by the dash-dot lines, is β_{2Hz} and $\alpha\beta_{8Hz}$ for the 2 and 8 Hz gratings respectively, where $\beta_{8Hz} > \beta_{2Hz}$ (that is, the inhibitory DC component is larger, as a percentage of the sinusoidal amplitude, for higher temporal frequencies). The zero level is shown by the dotted lines. The solid horizontal lines show the threshold; the response is the portion of the curves above threshold, showed by the filled areas. Because β_{8Hz} is larger than β_{2Hz}, the sinusoidal component at 8 Hz needs a higher amplitude to give the same F1 response as the 2 Hz grating; that is, to achieve equal F1 responses, we must have $\alpha > 1$, meaning that the 8 Hz grating has higher contrast

than the 2 Hz grating. The input in response to the two gratings is just the sum of the inputs in response to each grating; its mean is the sum of the two means, $\beta_{2Hz} + \alpha\beta_{8Hz}$. The resulting response shows an increased 8 Hz response and a decreased 2 Hz response, relative to the response to each grating alone; this occurs because $\beta_{8Hz} > \beta_{2Hz}$. See Fig. 12B-F for quantitative results (the values chosen for this illustration are for display only and do not correspond to those used to obtain the results in the simulations).

Figure 14

Normalized F1 response to eight single gratings at the preferred orientation, at their respective temporal frequencies ($\times$, dashed line), and to the same temporal frequencies for a superposition of the eight gratings ($+$, solid line). The contrast of the single gratings are 20% when alone, and 12.5% each when superposed. The normalized F1 response of the low temporal frequencies is higher for the single gratings than for the stimulus of eight superposed gratings, while the effect is reversed for high temporal frequencies, as found by Reid et al. (1992). Normalized F1 response is defined, as in Reid et al. (1992), as the F1 per unit contrast: it is F1/0.2 for the single sinusoids, and F1/(8×0.125) =F1 for the 8 sinusoids.

Figures

Figure 1:

A LGN input to simple cell

B Full circuit input to simple cell

Figure 2:

Figure 3:

Figure 4:

Figure 5:

Figure 6:

Figure 7:

Figure 8:

Figure 9:

Figure 10:

Figure 11:

A Single vs. double grating F1

B F1 with threshold

C F1 with 8Hz amplitude

D 2Hz F1 response difference

E 8Hz F1 response difference

F Valid parameter regime

Figure 12:

Figure 13:

Figure 14: